

Wait-Freedom with Advice

Carole Delporte-Gallet¹ Hugues Fauconnier¹

Eli Gafni² Petr Kuznetsov³

¹ University Paris Diderot

² UCLA

³TU Berlin/Deutsche Telekom Laboratories

Abstract

We motivate and propose a new way of thinking about failure detectors which allows us to define, quite surprisingly, what it means to solve a distributed task *wait-free using a failure detector*. In our model, the system is composed of *computation* processes that obtain inputs and are supposed to output in a finite number of steps and *synchronization* processes that are subject to failures and can query a failure detector. We assume that, under the condition that *correct* synchronization processes take sufficiently many steps, they provide the computation processes with enough *advice* to solve the given task wait-free: every computation process outputs in a finite number of its own steps, regardless of the behavior of other computation processes. Every task can thus be characterized by the *weakest* failure detector that allows for solving it, and we show that every such failure detector captures a form of set agreement. We then obtain a complete classification of tasks, including ones that evaded comprehensible characterization so far, such as renaming or weak symmetry breaking.

1 Introduction

What does it mean to solve a task? A distributed task for a set of processes can be seen as a function that maps an input vector to an output vector, one value per process. It is easy to reason about correctness of a task solution by matching the outputs to the inputs with respect to the task specification. When it comes to progress, however, it is getting less trivial.

On the surface, it is desirable to expect that the input vector is exactly matched by the output vector, i.e., every participating process obtains an output.¹ Unfortunately, in asynchronous or partially synchronous systems where relative processes' speeds are unbounded or very large, ensuring this property would require very long waiting. A more natural *wait-freedom* property requires that any participating process that takes sufficiently many steps obtains an output, “regardless of execution speeds of other processes” [20]. A wait-free task solution thus allows for treating the requirement “a given participant outputs” as a *liveness* property [2]: every execution has an extension in which the requirement is met. Naturally, wait-freedom assumes no notion of process *failures*: a process that does not take steps for a while in a given execution, always has a chance to wake up and take enough steps to output.

Failure detectors. Unfortunately, very few tasks can be solved wait-free in the basic read-write shared-memory model [14, 25, 21, 27, 5, 7]. The *failure detector* abstraction [9, 10] was proposed to circumvent these impossibilities. Intuitively, a failure detector provides each process with some (possibly incomplete and inaccurate) information about the current *failure pattern*, e.g., a list of processes predicted to take only finitely many steps in the current execution. The failure detector abstraction gives a language for capturing the weakest support from the system one may require in order to solve a given task. This gave many interesting insights on the nature of “wait-free unsolvable” tasks, starting from the celebrated result by Chandra et al. on the weakest failure detector for consensus [9].²

A solution of the task using a failure detector guarantees that every *correct* (a process that is predicted to take infinitely many steps by the failure pattern) eventually obtains an output. The progress of each process may thus depend on the behavior of other correct processes, and therefore failure detector-based algorithm cannot be wait-free. Consequently, since the failure pattern is introduced as a part of a run, we cannot treat individual progress as a liveness property anymore: a process is not allowed to take steps after it crashes.

Wait-freedom with advice. But can we think of a system where a “hard” task can be solved so that progress of a process does not depend on the execution speeds of other processes? A straightforward way to achieve this is to assume that the processes receive *advice* from an external oracle, and an immediate question is what is the weakest oracle that allows for solving a given task so that every participating process taking enough steps outputs.

In this paper, we use the language of failure detectors to determine the relative power of such external oracles. The oracle is represented as a set of *synchronization* processes equipped with a failure detector: each synchronization process can query its failure detector module to get hints about the failures of other synchronization processes. Thus, our system only considers failures of synchronization processes. As in the classical failure-detector literature [9], the assumptions about when and where failures of synchronization processes can occur are encapsulated in an *environment*, i.e., a set of allowed failure patterns. *Computation* processes (participants in a task solution) and synchronization processes communicate by reading and writing in the shared memory.

¹A process is considered participating if it takes at least one steps in the computation.

²Informally, \mathcal{D} is the weakest failure detector to solve a task T if it (1) solves T and (2) can be deduced from any failure detector that solves T .

Now what do we mean by solving a task with a failure detector? We require that, under the condition that the synchronization processes using their failure detector behave as predicted by the environment, every computation process taking enough steps must output.

It is easy to see that the classical failure-detector model [9] is a special case of our model where there is a bijective map between computation and synchronization processes, and a computation process stops taking steps after its synchronization counterpart does. Strictly speaking, when it comes to solving tasks, our framework demands from a failure detector more than the conventional failure detector model does. Indeed, in our framework, the failure of a synchronization process does not affect computation processes, and a failure detector is supposed to help computation processes output, as long as they take enough steps. In particular, we observe that the *weakest* failure detector to solve a task T in our framework is at least as strong as the weakest failure detector for T in the conventional model [9].

Ramifications. The idea of separating computation from synchronization is not new, e.g., it is used in the celebrated Paxos protocol [24] separating *proposers* from *acceptors* and *learners*. But applying it to distributed computing with failure detectors results in a surprisingly simple model, which we call *external failure detection (EFD)*, which resolves a number of long-standing puzzles.

The use of EFD enables a complete characterization of distributed tasks, based on the “amount of concurrency” they can stand. In the classical framework, we say that a task T can be solved k -concurrently if it guarantees that in every k -concurrent run every process taking sufficiently many steps eventually outputs [16]. Informally, a run is k -concurrent if at each moment of time, there are at most k participating processes without outputs. Now, in a system of n processes, each task T is associated with the largest k ($1 \leq k \leq n$) such that T can be solved k -concurrently.

We show that in EFD, a failure detector \mathcal{D} can be used to solve a task T with “concurrency level” at most k *if and only if* \mathcal{D} can be used to solve k -set agreement. More precisely, we show that, in every environment, i.e., for all assumptions on where and when failures of *synchronization* processes may occur, any failure detector that solves T is at least as strong as the anti- Ω_k failure detector [26, 28], denoted $\neg\Omega_k$. Then we describe an algorithm that uses $\neg\Omega_k$ to solve T (or any task that can be solved k -concurrently), in every environment.

Thus, any task is completely characterized through the “level of concurrency” its solution can tolerate. All tasks that can be solved k -concurrently but not $(k + 1)$ -concurrently (e.g., k -set agreement) are equivalent in the sense that they require exactly the same amount of information about failures (captured by $\neg\Omega_k$) to be solved in EFD. Note that this characterization covers *all* tasks, including “colored” ones that evaded any characterization so far [13, 18, 1].

Consider, for example, the task of (j, ℓ) -renaming in which j processes come from a large set of potential participants and choose new names in a smaller name space $1, \dots, \ell$, so that no two processes choose the same name. Surprisingly, in the conventional model, the renaming task itself can be formulated as a failure detector, so the question of the weakest failure detector for solving it results in a triviality. To avoid trivialities, additional assumptions on the scope of failure detectors are made [1].

In EFD, however, we immediately see that (j, j) -renaming (also called *strong* renaming) cannot be solved 2-concurrently and is thus equivalent to consensus.³ More generally, determining the weakest failure detector for (j, ℓ) -renaming boils down to determining the maximal k ($1 \leq k \leq j$) such that the task can be solved k -concurrently. We show finally that $(j, j + k - 1)$ -renaming can be solved k -concurrently, and, thus, using $\neg\Omega_k$.⁴

Another interesting corollary of our characterization is that if a failure detector solves k -set agreement among an arbitrary given subset of $k + 1$ processes, then it is strong enough to solve k -set agreement among *all* processes. This is a generalization of the recent result of Delporte et

³Note that all tasks can be solved 1-concurrently.

⁴For some values of j and k , however, the question of the maximal tolerated concurrency of $(j, j + k - 1)$ -renaming is still open [8].

al. [12] that any failure detector allowing for solving consensus (1-set agreement) among each two processes, also allows for solving consensus among all processes. Years of trying to show that the phenomenon demonstrated in [12] generalizes to all $k \geq 1$ in the conventional failure-detector model [9] bore no fruits.

One important feature inherited by our EFD framework from wait-free protocols is that it leverages *simulation-based* computing: processes can cooperate trying bring *all* participating processes to their outputs. Simulations were instrumental in establishing tight relations between seemingly different phenomena in *asynchronous* systems [5, 7, 15, 19, 13, 17, 16], and we extend this line of research below to failure-detector models.

Roadmap. The paper is organized as follows. First, we formally define our model and our new notion of task solvability with a failure detector. We then present a simple inductive proof of a generalization of [12] to any $k > 1$. Then we extend the generalization even further by presenting a complete characterization of decision tasks, based on the level of concurrency they can tolerate. Then we derive the weakest failure detector for strong renaming and wrap up with obligatory concluding remarks. Proofs are partially delegated to the optional Appendix.

2 The model of external failure detection

In this section, we propose a new definition of what it means to solve a task using a failure detector and relate it to the conventional definition of [9]. Parts of our model reuse elements of [9, 10, 18, 21].

2.1 Model for computation and synchronization

Our system is split in two parts. The *computation* part is made up of processes that get input values for the task they intend to solve and return output values. The *synchronization* part is made up of processes that use failure detectors to help processes of the computation part.

Processes. Formally, we consider a read-write shared-memory system which consists of m *C-processes*, $\Pi^C = \{p_1, \dots, p_m\}$, and n *S-processes*, $\Pi^S = \{q_1, \dots, q_n\}$. We allow n and m to be arbitrary natural numbers, but, as we shall see shortly, the only “interesting” case is when $n = m$.

Intuitively, the *C*-processes are responsible for computation. The *S*-processes are responsible for synchronization and may be equipped with a failure detector module [10] that gives hints about failures of other *S*-processes. The processes in $\Pi^C \cup \Pi^S$ communicate with each other via reading and writing in the shared memory.

Failure patterns and failure detectors. Since *C*-processes are assumed to be wait-free, we are only interested here in failures of *S*-processes. Hence a *failure pattern* F is a function from the time range $\mathbb{T} = \mathbb{N}$ to 2^{Π^S} , where $F(\tau)$ denotes the set of *S*-processes that have crashed by time τ . Once a process crashes, it does not recover, i.e., $\forall \tau : F(\tau) \subseteq F(\tau + 1)$. $\text{faulty}(F) = \cup_{\tau \in \mathbb{T}} F(\tau)$ is the set of faulty processes in F and $\text{correct}(F) = \Pi^S - \text{faulty}(F)$ is the set of correct processes in F .

A *failure detector history* H with range \mathcal{R} is a function from $\Pi^S \times \mathbb{T}$ to \mathcal{R} . $H(q_i, \tau)$ is interpreted as the value output by the failure detector module of *S*-process q_i at time τ . A *failure detector* \mathcal{D} with range $\mathcal{R}_{\mathcal{D}}$ is a function that maps each failure pattern to a (non-empty) set of failure detector histories with range $\mathcal{R}_{\mathcal{D}}$. $\mathcal{D}(F)$ denotes the set of possible failure detector histories permitted by \mathcal{D} for failure pattern F .

An *environment* \mathcal{E} is a set of failure patterns that describes a set of conditions on when and where failures might occur. For example \mathcal{E}_t is the environment that consists of all failure patterns F such that $\text{correct}(F) \geq n - t$. We assume that for every failure pattern in the environments we consider, at least one *S*-process is correct.

Algorithms and runs. A distributed algorithm \mathcal{A} using a failure detector \mathcal{D} consists of two collections of deterministic automata, $\mathcal{A}_1^C, \dots, \mathcal{A}_m^C$, one automaton for each C -process, and $\mathcal{A}_1^S, \dots, \mathcal{A}_n^S$, one automaton for each S -process. In a *step* of the algorithm, a process may read or write to a shared register, or (if it is a S -process) consult its failure-detector module.

A *state* of \mathcal{A} is defined as the state of each process (state of each process being identified with the state of its corresponding automaton) and each shared object in the system. An *initial state* I of \mathcal{A} specifies an initial state for every process and every shared object.

A *run* of \mathcal{A} using a failure detector \mathcal{D} in an environment \mathcal{E} is a tuple $R = \langle F, H, I, Sch, T \rangle$ where $F \in \mathcal{E}$ is a failure pattern, $H \in \mathcal{D}(F)$ is a failure detector history, I is an initial state of \mathcal{A} , Sch is an infinite *schedule*, i.e., a sequence of processes in $\Pi^C \cup \Pi^S$, T is a sequence of non-decreasing elements of \mathbb{T} . The k -th step of run R is a step of process $Sch[k]$ determined by the current state, the failure history H , $T[k]$ and the algorithm \mathcal{A} . If it is a step of a S -process, this process is alive ($Sch[k] \notin F(T[k])$) and the value of the failure detector for this step is given by $H(Sch[k], T[k])$.

Let $\inf^S(R)$ denote the set of processes in Π^S that appear infinitely often in Sch . Respectively, $\inf^C(R)$ denote the set of processes in Π^C that appear infinitely often in Sch . We say that a run $R = \langle F, H, I, Sch, T \rangle$ is *fair* if $\text{correct}(F)$ is equal to $\inf^S(R)$, and $\inf^C(R)$ is not empty. A *finite run* of \mathcal{A} is a “prefix” of a run $\langle F, H, I, Sch, T \rangle$ of \mathcal{A} , i.e., a tuple $\langle F, H, I, Sch', T' \rangle$ such that $|Sch'| = |T'|$, Sch' is a proper prefix of Sch , and T' is a proper prefix of T .

Tasks. We focus on a class of problems called *tasks* that are defined uniquely through inputs and outputs.

A *task* [21] is defined through a set \mathcal{I} of input vectors (one input value for each C -process), a set \mathcal{O} of output vectors (one output value for each C -process), and a total relation $\Delta : \mathcal{I} \mapsto 2^{\mathcal{O}}$ that associates each input vector with a set of possible output vectors. An input value equal to \perp denotes a *not participating* process and \perp output value denotes an *undecided* process.

A m -vector L' is a *prefix* of a m -vector L if L' contains at least one non- \perp item and for all i , $1 \leq i \leq m$, either $L'[i] = \perp$ or $L'[i] = L[i]$. A set \mathcal{L} of vectors is *prefix-closed* if for all L in \mathcal{L} every prefix of L is in \mathcal{L} .

We assume that each element of \mathcal{I} and \mathcal{O} contains at least one non- \perp item and also that the sets \mathcal{I} and \mathcal{O} are prefix-closed. Moreover, we only consider tasks that have finite sets of input vectors \mathcal{I} (this assumption is used in Section 4 when we categorize tasks based on the failure detectors needed to solve them).

We stipulate that if $(I, O) \in \Delta$, then (1) if, for some i , $I[i] = \perp$, then $O[i] = \perp$, (2) for each O' , prefix of O , $(I, O') \in \Delta$ and, (3) for each I' such that I is a prefix of I' , there exists some O' such that O is a prefix of O' and $(I', O') \in \Delta$.

For example, in the task of (U, k) -*agreement*, where $U \subseteq \Pi^C$, input and output vectors are m -vectors, such that $I[i] = \perp$ for all $p_i \notin U$, input values are in $\{\perp, 0, \dots, k\}$, output values are in $\{\perp, 0, \dots, k\}$, and for each input vector I and output vector O , $(I, O) \in \Delta$ if the set of non- \perp values in O is a subset of values in I of size at most k . (Π^C, k) -agreement is the conventional k -set agreement task [11] and $(\Pi^C, 1)$ -agreement is *consensus* [14].

2.2 Solving a task in the EFD framework

Now we are ready to define what does it mean to solve a task in the external failure detection framework.

Input vector and output vector of a run. First, we assume that each automaton \mathcal{A}_i^C (1) gets an input value input_i as part of its initial state, and (2) contains *decide* steps such that all the next steps of \mathcal{A}_i are null steps that do not affect the current state when they are executed and for each *decide* step is associated a decision value v_i .

The first step of each C -process is to write its input value to shared memory. A process that wrote its input value is called *participating*. If a C -process executes a *decide* step with decision

value v , we say that the process decides v or returns v .

Given a run R , the *input vector* for the run is the m -vector I such that $I(i) = \text{input}_i$ if p_i is a participating process and $I(i) = \perp$ if p_i is a not participating process. In the same way, the *output vector* of the run is the m -vector O such that $O(i) = v$ if p_i decides v in the run and $O(i) = \perp$ if p_i does not decide in the run.

Solving a task. We say that a run R with input vector I and output vector O *satisfies a task* $T = (\mathcal{I}, \mathcal{O}, \Delta)$ if (1) $(I, O) \in \Delta$ and (2) $O(i) = \perp$ only if p_i makes a finite number of steps ($p_i \notin \text{inf}^C(R)$).

An algorithm \mathcal{A} *EFD-solves a task* $T = (\mathcal{I}, \mathcal{O}, \Delta)$ *using a failure detector* \mathcal{D} *in an environment* \mathcal{E} (in the rest we simply say “solves”) if every fair run of \mathcal{A} satisfies T . If such an algorithm exists for task T , T is *solvable with failure detector* \mathcal{D} *in environment* \mathcal{E} . By extension, a failure detector \mathcal{D} *solves a task* T *in* \mathcal{E} if there is an algorithm \mathcal{A} that solves T using \mathcal{D} in \mathcal{E} .

Note that we expect the algorithm to guarantee output to every C -process that takes sufficiently many steps, regardless of where and when S -processes fail. The algorithm only expects that every correct S -process in the current failure pattern takes infinitely many steps.

Comparing failure detectors. Failure detector reduction is defined as usual: failure detector \mathcal{D}' is *weaker than failure detector* \mathcal{D} *in an environment* \mathcal{E} if S -processes can use \mathcal{D} to emulate \mathcal{D}' in \mathcal{E} . More precisely, the automata of the C -processes of the distributed *reduction algorithm* \mathcal{A} are automata with only null steps and the emulation of \mathcal{D}' using \mathcal{D} is made by maintaining, at each S -process q_i \mathcal{D}' -output $_i$ so that in any fair run with failure pattern F , the evolution of variables $\{\mathcal{D}'\text{-output}_i\}_{q_i \in \Pi^S}$ results in a history $H' \in \mathcal{D}'(F)$. We say that two failure detectors are *equivalent* in \mathcal{E} if each is weaker than the other in \mathcal{E} .

As in the original definition [9], if failure detector \mathcal{D}' is weaker than failure detector \mathcal{D} in environment \mathcal{E} , then every task solvable with \mathcal{D}' in \mathcal{E} can also be solved with \mathcal{D} in \mathcal{E} . Now \mathcal{D} is the *weakest failure detector* to solve a task T in \mathcal{E} if (i) \mathcal{D} solves T in \mathcal{E} and (ii) \mathcal{D} is weaker than any failure detector that solves T in \mathcal{E} . It is straightforward to extend the arguments of [22] to show that every task has a weakest failure detector.

k-concurrency. Consider the solvability of a task without the help of a failure detector. In this case the deterministic automata of the S -processes of the distributed algorithm \mathcal{A} are automata with only null steps. Such an algorithm will be called *restricted*.

It is clear that tasks that are solvable with a restricted algorithm are exactly tasks that are said *wait-free* solvable in the literature (e.g. in [20, 21]).

The notion of *k-concurrent* solvability, introduced in [16], is a weaker form of solvability: a task is solvable *k-concurrently* if it is solvable only when at most k C -processes concurrently invoke the task. More precisely, a run of a distributed algorithm is *k-concurrent* if it is fair and at each time there is at most k undecided participating C -processes. A task $T = (\mathcal{I}, \mathcal{O}, \Delta)$ is *k-concurrently* solvable if there is a restricted algorithm \mathcal{A} such that all *k-concurrent* runs R of \mathcal{A} satisfy T . Note that runs of \mathcal{A} in which the number of participating but not decided C -processes exceeds k at some point may not satisfy T .

A wait-free solvable task is *m-concurrently* solvable. Also, it is easy to show that:

Proposition 1 *Every task is 1-concurrently solvable.*

Restriction on the number of C-processes. Trivially, if a task T is solvable with a restricted algorithm then T is also solvable with any number of S -processes and any failure detector. Reciprocally, consider an algorithm \mathcal{A} solving a task T with a trivial failure detector⁵ in environment \mathcal{E}_{n-1} . If $n \geq m$ consider the following algorithm: each C -process p_i executes alternatively steps of $\mathcal{A}_{p_i}^C$ and steps of $\mathcal{A}_{q_i}^S$ and each S -process executes only null steps. It is easy to verify that in this

⁵A trivial failure detector always outputs \perp .

way we emulate runs of \mathcal{A} in the failure pattern in which at least all S -processes q_i with $i > m$ are crashed, and such runs satisfy task T . Hence we get:

Proposition 2 *If $n \geq m$, T is solvable in \mathcal{E}_{n-1} with a trivial failure detector if and only if T is solvable with a restricted algorithm.*

But if $n < m$, the S -processes may help solving the task even if they do not use their failure detection capacities. For example, with n S -processes we can implement a (Π^C, n) -set agreement in every environment. For this, each S -process waits until at least one C -process writes its input in shared memory, and then it writes this value to a shared variable V . Each C -process waits until V has been written and outputs the read value. As at least one S -process is correct, eventually V will be written and as there are n S -processes at most n values may be output. In this way the (Π^C, n) -set agreement is always solvable even without the help of any failure detector.

As we focus here on solvability where additional power of processes is only due to the failure detection, the only “interesting” scenario to consider is when the number of C -processes does not exceed the number of S -processes and more specifically the case where they are equal. *Therefore, in the following we assume that the number of C -processes is equal to the number of S -processes, we denote this number by n .*

2.3 Conventional solvability

More conventional models of computation in which there is no separation between the computation and the synchronization part may be considered as a special case of the generalized model presented here. In conventional models, each process $i \in \{1, \dots, n\}$ can be seen as running two parallel threads: p_i corresponding to the computational part and q_i corresponding to the synchronization part. Moreover failure patterns correspond: i is correct in conventional systems if and only if q_i is correct in our setting. But, since in our model, computation and synchronization are separate, it is possible that p_i makes only a finite number of steps even if q_i is correct or vice-versa. Then we define *personified runs* of a distributed algorithm as being runs R that are fair and such that p_i crashes if and only if q_i crashes at the same time (as a result, $\inf^C(R)$ is equal to $\inf^S(R)$). We say that algorithm \mathcal{A} solves *classically* task T with failure detector \mathcal{D} in environment \mathcal{E} if every personified run R of \mathcal{A} satisfied T .

This definition corresponds exactly to the notion of solvability in a conventional setting as can be found in the literature [9].

As the set of personified runs of a distributed algorithm is a subset of the fair runs, we have:

Proposition 3 *If a failure detector \mathcal{D} solves a task T in an environment \mathcal{E} then \mathcal{D} classically solves T in \mathcal{E} .*

Corollary 4 *If \mathcal{D} is the weakest failure detector to classically solve a task T in an environment \mathcal{E} , then \mathcal{D} is weaker than the weakest failure detector to solve T in \mathcal{E} .*

Note that the converse of Proposition 3 is not true. For example, consider the $(\{p_1, p_2\}, 1)$ -agreement task (consensus among p_1 and p_2). It is classically solvable in \mathcal{E}_2 (assuming at most 2 failures) with the failure detector \mathcal{D} that, for each S -process, outputs q_1 if q_1 is correct and outputs q_2 if q_1 is faulty. But this task is not solvable in \mathcal{E}_2 with this failure detector (intuitively, otherwise, if q_1 is crashed we would be able to solve consensus between p_1 and p_2 without a failure detector).

However for *colorless* tasks⁶ both notions of solvability coincide.

⁶Informally, in a solution of a colorless task [7], a process is free to adopt the input or the output value of any other participating process.

Proposition 5 *Let T be a colorless task, T is solvable with failure detector \mathcal{D} in environment \mathcal{E} if and only if T is classically solvable with \mathcal{D} in \mathcal{E} . The weakest failure detector to solve T in \mathcal{E} is the weakest failure detector to classically solve T in \mathcal{E} .*

Failure detectors for k -set agreement.

The failure detector $\neg\Omega_k$ [28] outputs, at every S -process and each time, a set of $(n - k)$ S -processes. $\neg\Omega_k$ guarantees that there is a time after which some correct S -process is never output:

$$\begin{aligned} \forall F \in \mathcal{E}, \forall H \in \neg\Omega_k(F), \exists q_i \in \text{correct}(F), \tau \in \mathbb{T}, \\ \forall \tau' > \tau, \forall q_j \in \text{correct}(F) : q_i \notin H(q_j, \tau'). \end{aligned}$$

$\neg\Omega_1$ is equivalent to Ω [9] that outputs a S -process such that eventually the same correct S -process is permanently output at all correct processes.

From [18], we know that in every environment \mathcal{E} , $\neg\Omega_k$ is the weakest failure detector to classically solve (Π^C, k) -set agreement in \mathcal{E} . As (Π^C, k) -set agreement is a colorless task, from Proposition 5 we obtain:

Proposition 6 *In every environment \mathcal{E} , $\neg\Omega_k$ is the weakest failure detector to solve (Π^C, k) -set agreement in \mathcal{E} .*

3 Solving a puzzle

Let U be a set of $k + 1$ C -processes. Consider a failure detector \mathcal{D} that *solves* k -set agreement among the processes in U . We show that \mathcal{D} actually solves k -set agreement among *all* n C -processes.

Theorem 7 *Let U be a set of $(k + 1)$ C -processes, for some $1 \leq k < n$. For every environment \mathcal{E} , if a failure detector \mathcal{D} solves (U, k) -set agreement in \mathcal{E} then \mathcal{D} solves (Π^C, k) -set agreement in \mathcal{E} .*

Proof sketch. Without loss of generality, assume that $U = \{p_1, \dots, p_{k+1}\}$. Let \mathcal{A} be a distributed algorithm that solves the (U, k) -set agreement in \mathcal{E} with \mathcal{D} .

Let U_x denote $\{p_1, \dots, p_x\}$, $x = k + 1, \dots, n$. We observe first that \mathcal{D} can be used to solve $(U_x, x - 1)$ -set agreement as follows. C -processes in $\{p_1, \dots, p_{k+1}\}$ and S -processes $\{q_1, \dots, q_n\}$ run \mathcal{A} to solve k -set agreement and return the value returned by the algorithm, and processes in $\{p_{k+2}, \dots, p_x\}$ simply return their own input values. In total, at most $x - 1$ distinct input values are returned. Let \mathcal{A}_x denote the resulting algorithm.

We proceed now by downward induction to show that for all $x = n$ down to k , \mathcal{D} solves (Π^C, x) -set agreement.

The base case is immediate: $\{p_1, \dots, p_n\}$ trivially solve (Π^C, n) -set agreement without any failure detector. Now suppose that \mathcal{D} solves (Π^C, x) -set agreement for $x \geq k + 1$. By Proposition 6, \mathcal{D} can be used to implement $\neg\Omega_x$.

Using the generic simulation technique presented in Appendix C.1, the C -processes, p_1, \dots, p_n , can use $\neg\Omega_x$ to simulate a run of the C -part of \mathcal{A}_x on p_1, \dots, p_x , so that at least one simulated process takes infinitely many steps.⁷ The S -part of \mathcal{A}_x is executed by S -processes. In the simulation, each simulating process proposes its input value as an input value in the first step for each simulated process in $\{p_1, \dots, p_x\}$ (this can be done, since (Π^C, x) -set agreement is a colorless task).

Suppose that the current run is fair, i.e., every correct S -process takes infinitely many steps. Therefore, we simulate a fair run of \mathcal{A}_x and thus eventually some simulated C -process in $\{p_1, \dots, p_x\}$

⁷ We could have used a “black-box” simulation of \mathcal{A}_x using (Π^C, x) -set agreement objects presented in [16]. To make the paper self-contained, we give a direct construction using $\neg\Omega_x$ in Appendix C.1.

decides on one of the input values of the C -processes. Once a simulator finds out that a simulated process decided, it returns the decided value. Thus, eventually, every correct simulator returns. Since all decided values come from a run of \mathcal{A}_x , at most $x - 1$ distinct input values can be decided. Hence, \mathcal{D} solves $(\Pi^C, x - 1)$ -set agreement. \square

Therefore, in our framework, we obtain a direct generalization of the fact that for a failure detector, it is as hard to solve consensus in a system of n processes as to solve consensus among each pair of processes [12]. In fact, the separation between C -processes and S -processes, implies a stronger result: solving k -set agreement among one given set of $(k + 1)$ processes is as hard (in the failure detector sense) as solving it among all n processes.

4 Generalizing the puzzle

We showed in the previous section that solving k -set agreement among any given set of $k + 1$ C -processes requires an amount of information about failures that is sufficient to solve k -set agreement among all n C -processes. We show below that this statement can be extended to any task T that cannot be solved $(k + 1)$ -concurrently. We present an explicit reduction algorithm that extracts $\neg\Omega_k$ from any failure detector that solves T . Conversely, we show that a task that is k -concurrently solvable can be solved with $\neg\Omega_k$ in any environment.

Finally, we derive a complete characterization of generic tasks : all tasks that can be solved k -concurrently but not $(k + 1)$ -concurrently are equivalent in the sense that they require the same information about failures to be solved ($\neg\Omega_k$).

4.1 Reduction to $\neg\Omega_k$

Let T be any task that cannot be solved $(k + 1)$ -concurrently. Let \mathcal{E} be any environment. We show that every failure detector \mathcal{D} that solves T in \mathcal{E} can be used to implement $\neg\Omega_k$ in \mathcal{E} as follows.

Let \mathcal{A} be the algorithm that solves T using \mathcal{D} in \mathcal{E} . Recall that \mathcal{A} consists of two parts: \mathcal{A}^C is run by the C -processes p_1, \dots, p_n and \mathcal{A}^S is run by the S -processes q_1, \dots, q_n .

First, we construct a restricted algorithm \mathcal{A}_{sim} . In \mathcal{A}_{sim} , C -processes p_1, \dots, p_n perform two parallel tasks. In the first task, C -processes take steps on behalf of \mathcal{A}^C . In the second task, they simulate a run of \mathcal{A}^S on S -processes using, instead of \mathcal{D} , a *directed acyclic graph* (DAG) G . The DAG G contains a sample of values output by \mathcal{D} in some run R of \mathcal{A} [9, 28]. In \mathcal{A}_{sim} , S -processes take null steps.

Informally, each run of \mathcal{A}_{sim} gives “turns” to the S -processes and if G provides enough information about failures to simulate the next step of a S -process q_j , the step of q_j appears in the simulated run of \mathcal{A} . To simulate steps of \mathcal{A}^S , C -processes employ BG-simulation [5, 7]. This simulation technique enables $k + 1$ processes called *simulators*, to simulate a run of any asynchronous n -processes protocol in which at least $(n - k)$ processes take infinitely many steps. Thus, if k or less participating C -processes take a finite number of steps, the resulting run of \mathcal{A}_{sim} gives infinitely many turns to at least $n - k$ S -processes.

Let F be the failure pattern of the run in which G was constructed. \mathcal{A}_{sim} guarantees that (1) every finite run of \mathcal{A}_{sim} simulates a finite run of \mathcal{A} , and (2) if every S -process that is correct in F receives infinitely many turns to take steps, then the simulated run of \mathcal{A} is fair, and (3) if k or less participating C -processes take only finitely many number of steps, then there are at most k S -processes that receive only finitely many turns to take steps in the simulation.

Second we construct a reduction algorithm. In such an algorithm C -processes take null steps. Our reduction algorithm consists of two components (both are run exclusively by the S -processes). In the first component, every S -process q_i queries \mathcal{D} , exchanges the returned values with other S -processes and maintains a DAG G_i . In the second component, each q_i locally simulates multiple $(k + 1)$ -concurrent runs of \mathcal{A}_{sim} using G_i , going over all combinations of inputs, exploring the

runs in the depth-first manner. The simulation continues as long as some simulated C -process does not decide in the produced run of \mathcal{A}_{sim} . Since T cannot be solved $(k + 1)$ -concurrently, there must be a $(k + 1)$ -concurrent run of \mathcal{A}_{sim} in which some participating C -process that takes infinitely many steps never decides. The only reason for a C -process not to decide in a run of \mathcal{A}_{sim} is that some correct S -process receives only finitely many turns in the simulation. But in the simulation, at least $(n - k)$ S -processes receive infinitely many turns. Thus, by outputting the identities of the $(n - k)$ S -processes that were last to receive turns in the current run we emulate the output of $\neg\Omega_k$: we output sets of $n - k$ S -processes that eventually never contain some correct process.

Theorem 8 *Let T be a task that cannot be solved $(k + 1)$ -concurrently. For every environment \mathcal{E} , for every failure detector \mathcal{D} that solves T in \mathcal{E} , $\neg\Omega_k$ is weaker than \mathcal{D} in \mathcal{E} .*

4.2 Solving a k -concurrent task with $\neg\Omega_k$

In this section, instead of $\neg\Omega_k$, we use an equivalent failure detector $\vec{\Omega}_k$ [28]. Basically, $\vec{\Omega}_k$ gives a k -vector of processes such that, eventually, at least one position of the vector stabilizes on the same correct process at all correct processes.

By definition if T is k -concurrently solvable, then there exists a restricted algorithm \mathcal{A} that k -concurrently solves T .

First, we define an abstract simulation technique that, with help of $\vec{\Omega}_k$, allows us to simulate, in a system of n C -processes, runs of any restricted input-less algorithm on k C -processes (the set of non- \perp input values is a singleton). Moreover, in this simulation, if ℓ simulators participate then at most $\min(k, \ell)$ processes take infinitely many steps in the simulated execution. Basically, to perform a step for a simulated C -process p_i , the C -processes and the S -processes execute an instance of a leader-based consensus algorithm [10], using the item i of $\vec{\Omega}_k$ as a leader. The property of $\vec{\Omega}_k$ ensures that for some i , infinitely many consensus instances terminate.

Second, we define a restricted algorithm \mathcal{B} for k C -processes that simulates a k -concurrent run of \mathcal{A} , using the BG-simulation techniques [6, 7]. Applying the abstract simulation technique to \mathcal{B} , we obtain an algorithm in which every run R simulates a run R_{sim} of \mathcal{A} such that: (1) R_{sim} contains only steps of participating processes of R , (2) the inputs of the participating processes are the same in R and R_{sim} , (2) R_{sim} is k -concurrent, and (3) every C -process that takes infinitely many steps in R takes also infinitely many steps in R_{sim} . So if T is k -concurrent solvable with \mathcal{A} , R_{sim} satisfies T , and, consequently, R satisfies T .

To sum up, we have constructed an algorithm that solves T with $\neg\Omega_k$: with the help of S -processes and $\neg\Omega_k$, p_1, \dots, p_n simulate C -processes p'_1, \dots, p'_k that, in turn, simulate C -processes p''_1, \dots, p''_n taking steps in a k -concurrent execution of algorithm \mathcal{A} .

Theorem 9 *Let T be any k -concurrently solvable task. For every environment \mathcal{E} , $\neg\Omega_k$ solves T in \mathcal{E} .*

4.3 Task hierarchy

From Theorems 8 and 9, we deduce:

Theorem 10 *Let T be a task that can be solved k -concurrently but not $(k + 1)$ -concurrently. In every environment \mathcal{E} , $\neg\Omega_k$ is the weakest failure detector to solve T in \mathcal{E} .*

As a corollary, all tasks that can be solved k -concurrently but not $(k + 1)$ -concurrently (e.g., k -set agreement) are equivalent in the sense that they require exactly the same amount of information about failures (captured by $\neg\Omega_k$).

5 Characterizing the task of strong renaming

To illustrate the utility of our framework, we consider the task of (j, ℓ) -renaming [3]. The task is defined on n ($n > j$) processes and assumes that in every run at most j processes participate (at least $n - j$ elements of each vector $I \in \mathcal{I}$ are \perp). As an output, every participant obtains a unique *name* in the range $\{1, \dots, \ell\}$ (every non- \perp element in each $O \in \mathcal{O}$ is a distinct value in $\{1, \dots, \ell\}$).

In this section, we first focus on (j, j) -renaming (also called *strong j -renaming*). Using Theorem 10, we show that the weakest failure detector for strong j -renaming is Ω (for each $1 < j < n$). In other words, strong renaming is equivalent to consensus.

Note that in strong 2-renaming at most 2 C -processes concurrently execute steps of the algorithm. So the impossibility to achieve strong 2-renaming is equivalent to the impossibility of solving strong 2-renaming 2-concurrently. By a simple reduction to the impossibility of wait-free 2-processes consensus, we show (Appendix D):

Lemma 11 *Strong 2-renaming cannot be solved 2-concurrently.*

By reducing to the impossibility of Lemma 11, we get a more general result:

Theorem 12 *For all $1 < j < n$, strong j -renaming cannot be solved 2-concurrently.*

Proposition 1, Theorem 10, and Theorem 12 imply:

Corollary 13 *For all j ($1 < j < n$), in every environment \mathcal{E} , Ω is the weakest failure detector for solving strong j -renaming in \mathcal{E} .*

In fact, there exists a generic algorithm (Appendix D.2) that, for all $k = 1, \dots, j$, solves $(j, j + k - 1)$ -renaming in all k -concurrent runs, and thus $(j, j + k - 1)$ -renaming can be solved using $\neg\Omega_k$. For some values of k and j , $(j, j + k - 1)$ -renaming can be shown to be impossible to solve $(k + 1)$ -concurrently, for others determining the maximal level of concurrency of $(j, j + k - 1)$ -renaming is still an open question [8].

6 Conclusion

This paper introduces a new model of distributed computing with failure detectors that allows processes to cooperate. A process in this model is able to advance the computation of other participating processes in the way used previously only in asynchronous simulations [5, 7, 15, 16], while using failure detectors to overcome asynchronous impossibilities. In our new framework, we derive a complete characterization of distributed tasks, based on their maximal “concurrency level”: class k ($1, \dots, n$) consists of tasks that can be solved at most k -concurrently, and all tasks in the class are equivalent to k -set agreement.

Our framework does not have to be tied to wait-freedom. We can think of its generalization to any progress condition on computation processes encapsulated, e.g., in an *adversary* [13]. Therefore, we can pose questions of the kind: what is the weakest failure detector to solve a task T in the presence of an adversary \mathcal{A} ? This gives another dimension to the questions explored in this paper.

References

- [1] Y. Afek and I. Nir. Failure detectors in loosely named systems. In *PODC*, pages 65–74. ACM Press, 2008.
- [2] B. Alpern and F. B. Schneider. Defining liveness. *Inf. Process. Lett.*, 21(4):181–185, Oct. 1985.
- [3] H. Attiya, A. Bar-Noy, D. Dolev, D. Peleg, and R. Reischuk. Renaming in an asynchronous environment. *Journal of the ACM*, 37(3):524–548, 1990.
- [4] H. Attiya and J. Welch. *Distributed Computing. Fundamentals, Simulations, and Advanced Topics*. John Wiley & Sons, 2004.
- [5] E. Borowsky and E. Gafni. Generalized FLP impossibility result for t -resilient asynchronous computations. In *STOC*, pages 91–100. ACM Press, 1993.
- [6] E. Borowsky and E. Gafni. Immediate atomic snapshots and fast renaming. In *PODC*, pages 41–51. ACM Press, 1993.
- [7] E. Borowsky, E. Gafni, N. A. Lynch, and S. Rajsbaum. The BG distributed simulation algorithm. *Distributed Computing*, 14(3):127–146, 2001.
- [8] A. Castañeda and S. Rajsbaum. New combinatorial topology bounds for renaming: the lower bound. *Distributed Computing*, 22(5-6):287–301, 2010.
- [9] T. D. Chandra, V. Hadzilacos, and S. Toueg. The weakest failure detector for solving consensus. *J. ACM*, 43(4):685–722, July 1996.
- [10] T. D. Chandra and S. Toueg. Unreliable failure detectors for reliable distributed systems. *J. ACM*, 43(2):225–267, Mar. 1996.
- [11] S. Chaudhuri. More choices allow more faults: Set consensus problems in totally asynchronous systems. *Information and Computation*, 105(1):132–158, 1993.
- [12] C. Delporte-Gallet, H. Fauconnier, and R. Guerraoui. Tight failure detection bounds on atomic object implementations. *J. ACM*, 57(4), 2010.
- [13] C. Delporte-Gallet, H. Fauconnier, R. Guerraoui, and A. Tielmann. The disagreement power of an adversary. *Distributed Computing*, 24(3-4):137–147, 2011.
- [14] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, Apr. 1985.
- [15] E. Gafni. The extended BG-simulation and the characterization of t -resiliency. In *STOC*, pages 85–92. ACM Press, 2009.
- [16] E. Gafni and R. Guerraoui. Generalized universality. In *CONCUR 2011 - Concurrency Theory - 22nd International Conference*, pages 17–27, 2011. Full version: <http://infoscience.epfl.ch/record/150307>.
- [17] E. Gafni and P. Kuznetsov. Turning adversaries into friends: Simplified, made constructive, and extended. In *OPODIS*, pages 380–394, 2010.
- [18] E. Gafni and P. Kuznetsov. On set consensus numbers. *Distributed Computing*, 24(3-4):149–163, 2011.

- [19] E. Gafni and P. Kuznetsov. Relating L -Resilience and Wait-Freedom via Hitting Sets. In *ICDCN*, pages 191–202, 2011. Full version: <http://arxiv.org/abs/1004.4701>.
- [20] M. Herlihy. Wait-free synchronization. *ACM Trans. Prog. Lang. Syst.*, 13(1):123–149, Jan. 1991.
- [21] M. Herlihy and N. Shavit. The topological structure of asynchronous computability. *J. ACM*, 46(2):858–923, 1999.
- [22] P. Jayanti and S. Toueg. Every problem has a weakest failure detector. In *PODC*, pages 75–84, 2008.
- [23] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, July 1978.
- [24] L. Lamport. The Part-Time parliament. *ACM Transactions on Computer Systems*, 16(2):133–169, May 1998.
- [25] M. Loui and H. Abu-Amara. Memory requirements for agreement among unreliable asynchronous processes. *Advances in Computing Research*, 4:163–183, 1987.
- [26] M. Raynal. K -anti-Omega, August 2007. Rump session at PODC 2007.
- [27] M. Saks and F. Zaharoglou. Wait-free k -set agreement is impossible: The topology of public knowledge. *SIAM J. on Computing*, 29:1449–1483, 2000.
- [28] P. Zieliński. Anti-Omega: the weakest failure detector for set agreement. *Distributed Computing*, 22(5-6):335–348, 2010.

A Proof for 1-concurrent solvable (Section 2.2)

Proposition 1 *Every task is 1-concurrent solvable.*

Proof. Each C -process p_i executes the following code (1) writes its input, (2) reads the other inputs already written getting a vector I such that $I[i] = \perp$, and (3) reads all the other outputs already written getting a vector O . If O is only composed with \perp then p_i is the first process, it chooses an output according to its input and Δ . Otherwise let I' obtained from I by replacing the i -th item with the input value of p_i . By definition of tasks, if $(I, O) \in \Delta$, there exists a vector O' obtained from O by replacing the i item by a non \perp value such that $(I', O') \in \Delta$. Then p_i decides and outputs value $O'[i]$. Let R be a 1-concurrent run, by an easy induction on the number of participating processes we prove that R satisfies T . \square

B Proof for the reduction to $\neg\Omega_k$ (Section 4.1)

The algorithm sketched in Figure 1 describes the steps to be taken by S -processes q_1, \dots, q_n to emulate $\neg\Omega_k$. First we describe the asynchronous algorithm \mathcal{A}_{sim} used by the C -processes to simulate runs of \mathcal{A} , given a sample of the output of \mathcal{D} . Then we describe how the S -processes use multiple simulated runs of \mathcal{A}_{sim} to emulate the output of $\neg\Omega_k$.

Asynchronous simulation of \mathcal{A} . Following the technique of Chandra et al. [9], we represent a sample of the failure-detector output in the form of a directed acyclic graph (DAG). The DAG is constructed by the S -processes by periodically querying \mathcal{D} and collecting the output values: every vertex of the DAG has the form $[q_i, d, k]$ which conveys that the k -th query of \mathcal{D} performed by process q_i returned value d . An edge between vertexes $[q_i, d, k]$ and $[q_j, d', k']$ conveys that the k -th query of \mathcal{D} performed by q_i *causally precedes* [23] the k' -th query of \mathcal{D} performed by q_j .

As in [28, 18], any such DAG G can be used to construct a *restricted* algorithm \mathcal{A}_{sim} .

In \mathcal{A}_{sim} the C -processes p_1, \dots, p_n simulates runs of \mathcal{A} . The C -processes obtain input values for T and perform two parallel tasks. First, the C -processes take steps on behalf of \mathcal{A}^C . Second, they use BG-simulation [6, 7] to simulate a run of \mathcal{A}^S on q_1, \dots, q_n . But to simulate step of S -process instead of \mathcal{D} they use the information provided by G . More precisely, in the simulation, every S -process q_i takes steps as prescribed by \mathcal{A}^S , except that when q_i is about to query \mathcal{D} , it chooses the next vertex $[q_i, d, k]$ causally succeeding the latest simulated steps of \mathcal{A}^S of all S -processes seen by q_i so far. If G was constructed in a run of \mathcal{A} with failure pattern F , it is guaranteed that (1) every finite run simulated by \mathcal{A}_{sim} is a run of \mathcal{A} with failure pattern F , and (2) if the run of \mathcal{A}_{sim} contains infinitely many simulated steps of processes in $correct(F)$ then the simulated run is a fair run of \mathcal{A}^S with failure pattern F [28, 18].

\mathcal{A}^S does not have inputs. Therefore, the simulation tries to promote all n S -processes (but succeed to take step for a S -process q_i if there is enough value for q_i in G).

If the simulated run of \mathcal{A} generates an output value for p_i , p_i outputs this value and leaves the computation. Note that since T cannot be solved $(k+1)$ -concurrently, and all runs of \mathcal{A} are safe, there must be a $(k+1)$ -concurrent (simulated) run of \mathcal{A} in which some participating process takes infinitely many steps without outputting a value.

Extracting $\neg\Omega_k$. Now to derive $\neg\Omega_k$, each S -process in $i \in \{1, \dots, k\}$ collects the output of \mathcal{D} in G and simulates locally multiple $(k+1)$ -concurrent runs of \mathcal{A}_{sim} . The runs are simulated in the *corridor-based* depth-first manner [18] that works as follows.

We assume a total order on the subsets $P \subseteq \Pi^C$ so that if $P \subset P'$ then P precedes P' in the order. Each initial state I and each *schedule* σ , a sequence specifying the order in which

```

1  for all  $I_0$ , input vectors of  $T$  (in some order) do
    /* All possible inputs for  $p_1, \dots, p_n$  */
2  for all  $\pi_0$ , permutations of  $p_1, \dots, p_n$  (in some order) do
    /* All possible ‘‘arrival orders’’ */
3       $P_0 :=$  the set of first  $k + 1$   $C$ -processes in  $\pi_0$ 
4       $explore(I_0, \perp, P_0, \pi_0)$ 

5  function  $explore(I, \sigma, P, \pi)$ 
6       $\neg\Omega_k\text{-output}_i := n - k$   $S$ -processes that appear the latest in  $\alpha_i(I, \sigma)$ 
        (any  $n - k$   $S$ -processes if not possible)
7      if  $\exists q_j \in \Pi^S: \forall \sigma' \in \text{dom}(\alpha_i), \exists \sigma'', \text{ a prefix of } \sigma': \alpha_j(I, \sigma'') \text{ is deciding}$  then
        /* If all schedules explored so far were found deciding by  $q_j$  */
8          adopt  $q_j$ 's simulation
9      else
10          $N :=$  the set of undecided processes in  $(I, \sigma)$ 
11         for all  $p_j \in P - N$  do /* For each decided process in  $P$  */
12              $P := P - \{p_j\}$ 
13              $P := P \cup \{\text{the first process in } \pi \text{ that does not appear in } \sigma\}$ 
            /* Replace  $p_j$  with the next non-participant in  $\pi$  */
14         for all  $P' \subseteq P$  (in some order consistent with  $\subseteq$ ) do
            /* For all ‘‘sub-corridors’’ */
15             for all  $p_j \in P'$  (in  $\pi$ ) do
16                  $explore(I, \sigma \cdot p_j, P', \pi)$ 

```

Figure 1: Deriving $\neg\Omega_k$: code for each S -process q_i .

p_1, \dots, p_n take steps of \mathcal{A}_{sim} , determine a unique run of \mathcal{A}_{sim}^S simulated at process q_i , denoted $\alpha_i(I, \sigma)$.

For a given input vector I and a given permutation π of p_1, \dots, p_n , that describes the order in which the C -processes ‘‘arrive’’ at the computation. Initially, we select a set P of the first $k + 1$ processes in π as the participating set. Subsets $P' \subseteq P$ are then explored as ‘‘corridors’’ (line 16), in the deterministic order, from the narrowest (solo) corridors to wider and wider ones. Recursively, we go through simulating all runs in which only C -processes in P' take steps. In the course of simulation, if a participating C -process p_j decides, we replace it with a process that has not yet taken steps in the current computation (line 13). Since we only replace a decided process with a ‘‘fresh’’ non-participant, the participating set keeps the size of $k + 1$ or less processes. This procedure is repeated until every C -process decides. Thus, every simulated run is $(k + 1)$ -concurrent. Once the exploration of the current corridor is complete (the call of $explore$ in line 16 returns), we proceed to the next corridor, etc.

If, at some point, q_i finds out that another S -process q_j made more progress in the simulation (simulated more runs than q_i), then q_i ‘‘adopts’’ the simulation of q_j (line 8) by adopting q_j 's version of the DAG and the map α_j [18].

The output of $\neg\Omega_k$ is evaluated as the set of the ids of the latest $n - k$ processes in q_1, \dots, q_n that appear in the run of \mathcal{A}_{sim}^S in the currently simulated run of \mathcal{A}_{sim} (line 6).

Recall that T cannot be solved $(k + 1)$ -concurrently and thus there must exist a $(k + 1)$ -concurrent run of \mathcal{A}_{sim} in which some participating live process never decides. Since the only reason for the run of \mathcal{A}_{sim} not to decide is the absence of some correct process in the simulated k -resilient run of \mathcal{A}_{sim}^S , and the emulated output eventually never contains some correct process— $\neg\Omega_k$ is emulated. Thus:

Theorem 8 *Let T be a task that cannot be solved $(k + 1)$ -concurrently. For every environment \mathcal{E} , for every failure detector \mathcal{D} that solves T in \mathcal{E} , $\neg\Omega_k$ is weaker than \mathcal{D} in \mathcal{E} .*

Proof sketch. Our reduction algorithm works as follows. Every S -process q_i runs two parallel

tasks. First, it periodically queries its module of \mathcal{D} and maintains its directed acyclic graph G_i , as in [9, 18]. Second, it uses G_i to locally simulate multiple runs of \mathcal{A}_{sim} and emulates the output of $\neg\Omega_k$. Consider any run of the reduction algorithm. Let F be the failure pattern of that run.

First we observe that every simulated run of \mathcal{A}_{sim} is $(k+1)$ -concurrent. Indeed, initially, exactly $(k+1)$ C -processes participate and a new participant joins only after some participating C -process decides and departs.

Then we show that the correct S -processes eventually perform the same infinite sequence of recursive invocations of *explore*: *explore*(I, \perp, P_0, π) invokes *explore*(I, σ_1, P_1, π), which in turn invokes *explore*(I, σ_2, P_2, π), etc. (line 14). Indeed, all S -processes perform the simulations in the same order and since, the task is not $(k+1)$ -concurrently solvable, there must be a never deciding $(k+1)$ -concurrent run of \mathcal{A}_{sim} . Since all these P_ℓ are non-empty, there exists ℓ^* and P^* such that $\forall \ell \geq \ell^*, P_\ell = P^*$. Since we proceed from narrower corridors to wider ones, P^* is the set of live C -processes that never decide in the “first” never deciding $(k+1)$ -concurrent simulated run with a schedule σ^* .

Now we observe that all simulated runs eventually always extend a prefix $\bar{\sigma}^*$ of σ^* in which some simulated processes not in P^* already took all their steps in σ^* . Moreover, there is a time after all explored extensions of $\bar{\sigma}^*$ only contain steps of processes in P^* . By the properties of BG-simulation [5, 7], every S -process that appears only finitely often in the run of \mathcal{A}_{sim} simulated by σ^* (we called these processes *blocked* by σ^*) eventually never appears in all simulated run of \mathcal{A} . Let U be the set of S -processes blocked by σ^* . Since the run of \mathcal{A}_{sim} simulated by σ^* is $(k+1)$ -concurrent, processes in U eventually never appear among the last $n-k$ processes in $\alpha(I, \sigma)$ (line 6).

Now we observe that U must contain a correct (in F) S -process. If it is not the case, i.e., U doesn't contain a correct S -process, then the simulated run of \mathcal{A} is fair and thus the simulated run of \mathcal{A} must be deciding.

Thus, eventually some correct S -processes never appear in $\neg\Omega_k\text{-output}_i$ at every correct S -process $q_i \rightarrow \neg\Omega_k$ is emulated. \square

C Proof for solving a k -concurrent solvable task with $\neg\Omega_k$ (Section 2.2)

This section presents a distributed algorithm that uses $\neg\Omega_k$ to solve, in any environment, any task that can be solved k -concurrently. The result could have been obtained from the simulation of k -concurrency using (black-box) k -set agreement objects [16]. But for the sake of self-containment, we present a (simpler) direct construction of a k -concurrent run using $\neg\Omega_k$.⁸

First we describe an abstract simulation technique that uses $\vec{\Omega}_k$ (equivalent to $\neg\Omega_k$ [28]) to simulate, in a system of n C -processes, a run of an arbitrary asynchronous algorithm \mathcal{B} on k C -processes.

Then we apply this technique to show that, in every environment, we can use $\vec{\Omega}_k$ to simulate a run R_{sim} of any given n C -processes protocol \mathcal{A} . If R is the current run, we have the following properties: (1) R_{sim} only contains steps of participating processes of R , (2) R_{sim} is k -concurrent, and (3) every participating C -process of R that takes infinitely many steps is given enough steps in R_{sim} to decide.

C.1 Simulating k codes using $\neg\Omega_k$

Suppose we are given a read-write algorithm \mathcal{B} on k C -processes, p'_1, \dots, p'_k . Assuming that $\vec{\Omega}_k$ is available, the algorithm in Figure 2 describes how n *simulators*, C -processes p_1, \dots, p_n can

⁸The construction is similar to the one presented in [18] for the actively k -resilient case.

simulate an infinite run of \mathcal{B} .

The simulation is similar in spirit to BG-simulation [5, 7]. Every simulator p_i first registers its participation in the shared memory and then tries to advance simulated C -processes $p'_1, \dots, p'_{\min(k,m)}$, where m is the number of simulators that p_i has witnessed participating.

To simulate a step of p'_j , simulators agree on the view of the C -process after performing the step. However, instead of the BG-agreement protocol of [5, 7], we use here a *leader-based* consensus algorithm [9]. In the algorithm, a process periodically (in every round r of computation), queries the current leader to get an estimate of the decision.

Since in our algorithm both C -processes and S -processes can be elected leaders, we modify the algorithm of [9] as follows. When a process wants to get an estimate of the decision (say in round r), it publishes a query ($query, est', r$) in the shared memory (proposing its current estimate est'), waits until the current leader publishes a response (est, r), and adopts the estimate. For simplicity, we assume that every process (be it a C -process or a S -process) periodically scans the memory to find new queries of the kind ($query, est', r$) and responds to them by publishing one of the proposed estimates. Furthermore, we assume that each S -process periodically updates the shared array $\vec{\Omega}_k-S[1, \dots, k]$ with the output of its module of $\vec{\Omega}_k$. Recall that eventually some position $\vec{\Omega}_k-S[j]$ ($j \in \{1, \dots, k\}$) stabilizes on the identity of some correct S -process.

The resulting algorithm terminates under the condition that all C -processes eventually agree on the same correct leader. The instance of the consensus algorithm used to simulate ℓ -th step of C -process p'_j is denoted by $cons_{j,\ell}$.

The rule to elect the leader is the following. As long as the number of participating simulators is k or less, the participating simulator with the j -th smallest identity acts as a leader for simulating steps of p'_j . When the number of participating simulators exceeds k , the leader for simulating steps of p'_j is given by $\vec{\Omega}_k-S[j]$.

In both cases, at least one simulated C -process is eventually associated with the same correct leader. Thus, at least one simulated C -process makes progress in the simulation.

The algorithm also assumes that a simulator p_i may decide to leave the simulation if the simulated run produced a desired output (line 28). We use this option in the next section.

Theorem 14 *In every environment, the protocol in Figure 2 simulates an infinite run of any k -processes algorithm \mathcal{B} (as long as there is at least one not decided participating simulated process). Moreover, if ℓ simulators participate, i.e., $|pars| = \ell$, then at most $\min(k, \ell)$ processes participate in the simulated run.*

Proof. Consider an infinite run of the algorithm. Since every next state of each simulated process p'_j is decided using a consensus algorithm, every simulator observes exactly the same evolution of states for every simulated process. Thus, the simulated schedule indeed belongs to a run of \mathcal{B} .

Now consider the construction of variables $Leader_1, \dots, Leader_k$ used by the consensus algorithms $cons_{1,\ell}, \dots, cons_{k,\ell}$ (lines 31-36). Let ℓ be the number of participating simulators.

If $\ell \leq k$, the simulator with the j -th smallest identity in $pars$ is assigned to be the leader of exactly one simulated process p'_j . Since at least one simulator is correct, there exists p'_j ($j = 1, \dots, |pars|$) such that all instances $cons_{j,\ell_j}$ using $Leader_j$ eventually terminate. Thus, p'_j accepts infinitely many steps in the simulated run.

If $\ell > k$, at least one $Leader_j$ ($j = 1, \dots, k$) eventually stabilizes on some correct process identity, as guaranteed by the properties of $\vec{\Omega}_k$. Again, p'_j takes infinitely many steps in the simulated run.

In both cases, at most $\min(\ell, k)$ simulated processes appear in the produced run of \mathcal{B} , and at least one simulated process takes infinitely many steps. \square

Shared variables:

$R_j, j = 1, \dots, m$, initially \perp
 $V_j, j = 1, \dots, k$, initially the initial state of p'_j
 $\vec{\Omega}_{k-S}[j], j = 1, \dots, n$, initially q_1

Local variables:

$Leader_j, j = 1, \dots, k$, initially p_1
 $\ell_j, j = 1, \dots, k$, initially 1
 $v_j, j = 1, \dots, k$, initially \perp

Task 1:

```

17  $R_i := 1$ 
18  $undecided := true$ 
19 for  $j = 1, \dots, k$  do  $v_j := \{V_1, \dots, V_k\}$ 
20 while  $undecided$  do
21   for  $j = 1, \dots, \min(|pars|, k)$  do
22     perform one more step of  $cons_{j, \ell_j}(v_j)$  using  $Leader_j$  as a leader
23     if  $cons_{j, \ell_j}(v_j)$  returns  $v$  then { The next state of  $p'_j$  is decided }
24        $V_j := v$  { Adopt the decided state of  $p'_j$  }
25       simulate the next step of  $p'_j$  in  $\mathcal{B}$ 
26       if  $v$  allows  $p_i$  to decide then { The simulator can depart }
27          $undecided := false$ 
28          $R_i := \perp$ 
29          $v_j := \{V_1, \dots, V_k\}$  { Evaluate the next state of  $p'_j$  }
30          $\ell_j := \ell_j + 1$ 

```

Task 2:

```

31 while  $true$  do
32    $pars := \{p_j, R_j \neq \perp\}$ 
33   if  $|pars| \leq k$  then
34     for  $j = 1, \dots, |pars|$  do  $Leader_j :=$  the  $j$ -th smallest process in  $pars$ 
35   else
36     for  $j = 1, \dots, k$  do  $Leader_j := \vec{\Omega}_{k-S}[j]$ 

```

Figure 2: Simulating k codes using vector- Ω_k : the program code for simulator p_i

C.2 Solving a k -concurrent task with $\neg\Omega_k$

Theorem 9 *Let T be any k -concurrently solvable task. In every environment \mathcal{E} , $\neg\Omega_k$ solves T in \mathcal{E} .*

Proof. Let \mathcal{A} be the algorithm that solves T k -concurrently. We simply employ the simulation protocol in Figure 2 (Theorem 14), and suppose that the simulated algorithm \mathcal{B} is Extended BG-simulation [15] for \mathcal{A} . More precisely, \mathcal{B} simulates with k C -processes the algorithm \mathcal{A} with n C -processes.

Thus, the double simulation is built as follows. Every process p_i writes its input value of T to the shared memory and starts the simulation of k processes p'_1, \dots, p'_k using the algorithm in Figure 2. The simulated processes p'_1, \dots, p'_k run, in turn, BG-simulation of \mathcal{A} on n processes p''_1, \dots, p''_n .

Each simulated process p''_j is simulated only if the corresponding p_j has written its input of T in the shared memory and p''_j has not yet obtained an output in the simulated run. Moreover, to make sure that the simulation indeed produces a k -concurrent run, at any point of the simulation, each simulator in $p'_j \in \{p'_1, \dots, p'_k\}$ tries to advance the participating and not yet decided process with the smallest id. If the currently simulated process is found blocked [5, 7], i.e., the process cannot advance because another simulator started simulating a step of it but has not yet finished, p'_j proceeds to the next smallest undecided participating process in $\{p''_1, \dots, p''_n\}$. Since there are at most k simulators, at most $k - 1$ undecided participating processes can be found blocked and thus there are at most k undecided participating processes at a time—the resulting simulated run is k -concurrent.

When p''_j obtains an output, the corresponding simulator p_j considers itself “decided” (line 26), writes \perp in R_i (line 28) and departs.

If p'_i cannot make progress because each code it tries to simulate is blocked and there are no more codes to add, it “aborts” all blocked agreements [15] and resumes the simulation. Since, at each point of time, the number of simulated codes does not get below the number of simulators that take steps, the simulation keeps making progress.

Thus, as long as ℓ processes $\{p_{j_1}, \dots, p_{j_\ell}\}$ participate, only $\min(k, \ell)$ processes in $\{p''_1, \dots, p''_n\}$ take steps, which results in a k -concurrent simulated run of \mathcal{A} . Every process p''_j that takes steps eventually decides in a k -concurrent run of \mathcal{A} and the corresponding simulator p_j departs. As soon as the decided process p_i departs by writing \perp to R_i , we have one simulator p_i and one simulated process p''_i less. Therefore, as long as there is a simulator taking steps and the run is fair, the simulated run makes progress, i.e., more and more participants decide. Thus, we obtain an algorithm that, in every environment, solves T . \square

D Proof for characterizing the task of renaming (Section 5)

To illustrate the utility of our framework, we consider the task of (j, ℓ) -renaming [3]. The task is defined on n ($n > j$) processes and assumes that in every run at most j processes participate (at least $n - j$ elements of each vector $I \in \mathcal{I}$ are \perp). As an output, every participant obtains a unique *name* in the range $\{1, \dots, \ell\}$ (every non- \perp element in each $O \in \mathcal{O}$ is a distinct value in $\{1, \dots, \ell\}$).

We show first that (j, j) -renaming (also called *strong j -renaming*) is not 2-concurrently solvable. Then we present a generic algorithm that, for all $k = 1, \dots, j$, solves $(j, j + k - 1)$ -renaming in all k -concurrent run, and thus $(j, j + k - 1)$ -renaming can be solved (in IFD) using $\neg\Omega_k$.

D.1 Impossibility of 2-concurrent strong 2-renaming

Lemma 11 *Strong 2-renaming cannot be solved 2-concurrently.*

```

Shared variables:  $R_\ell$ ,  $\ell = 1, \dots, n$ , initially  $\perp$ 
37  $R_i := 1$  /* register participation */
38 repeat
39    $S := \{p_\ell \mid R_\ell \neq \perp\}$  /*get the current participating set */
40    $S' := \{p_\ell \mid R_\ell = 1\}$  /*get the set of not yet decided participants */
41    $min_1 := \min(S')$ 
42   if ( $|S'| = 1$ ) then  $min_2 := min_1$  else  $min_2 := \min(S' - \min(S'))$ 
43   if ( $|S| = j$  and ( $p_i = min_1$  or  $p_i = min_2$ )) or ( $|S| = j - 1$  and  $p_i = min_1$ ) then
44     take one more step of  $\mathcal{A}$  /*if among two not decided with smallest ids */
45 until decided
46  $R_i := 0$ 
47 return the name decided in  $\mathcal{A}$ 

```

Figure 3: A 1-resilient strong j -renaming algorithm: code for each C -process p_i .

Proof. We start with showing that for the special case of $j = 2$, strong renaming cannot be solved 2-concurrently. Suppose, by contradiction, that there exists a (restricted) algorithm \mathcal{A} that solves $(2, 2)$ -renaming 2-concurrently. Since we assumed $j < n$, we have at least 3 processes in the system. By the pigeon-hole principle, there exist two processes that decide on the same name $v \in \{1, 2\}$ in their solo runs of \mathcal{A} . Without loss of generality, let these processes be p_1 and p_2 and let v be 1.

Now p_1 and p_2 can wait-free solve 2-processes consensus as follows. Each process publishes its input and then runs \mathcal{A} until it obtains a name. If the name is 1, the process decides on its input, otherwise it decides on the input of the other process. Since a process in $\{p_1, p_2\}$ obtains 1 as a name in a solo run of \mathcal{A} , if 1 is not obtained, then the other process participates in the run of \mathcal{A} and, thus, has previously written its input. Therefore, every decided value was previously proposed. Since every obtained name is distinct, the two processes cannot decide on different values. This conclude the proof that strong 2-renaming cannot be 2-concurrently solvable. \square

Theorem 12 *For all $1 < j < n$, strong j -renaming cannot be solved 2-concurrently.*

Proof. By Lemma 11, we have already the result for $j = 2$. Suppose, by contradiction, that for some $2 < j < n$, there exists an (restricted) algorithm \mathcal{A} solving strong j -renaming 2-concurrently. As we deal here with 2-concurrent solvability, we are only interested by the C -processes and their algorithms. We use \mathcal{A} to solve strong j -renaming in all 1-resilient runs, i.e., runs in which at least $j - 1$ C -processes participate and take infinitely many steps. Recall that at most j C -processes participate in every run, so either $j - 1$ or j processes take infinitely many steps. In the algorithm (Figure 3), every process registers its participation (line 37) and then periodically checks the current set of participants (line 39). If it finds out that it is among 2 processes with the smallest identities among j participating but not yet processes (line 43), then it starts taking steps \mathcal{A} until the algorithm provides p_i with a new name. Then p_i declares that it has decided (line 46) and departs.

Note that the resulting run of \mathcal{A} is 2-concurrent: either the participating set is of size $j - 1$ and only the not yet decided participant with the smallest identity is allowed to take steps of \mathcal{A} solo, or exactly j processes participate and the two not yet decided processes with the smallest identity are allowed to take step concurrently.

Now we observe that the run of \mathcal{A} continues as long as there is at least one not yet decided participant that take steps. Indeed, either the participating set is of size $j - 1$ and every participant takes an infinity number of steps (including the not yet decided one with the smallest identity) or exactly j C -processes participate and at least one of the not yet decided processes with the two smallest identity takes an infinity number of steps. Thus, every C -process that keeps taking

Shared variables:

$R_\ell, \ell = 1, \dots, n$, initially \perp

```

48   $s := 1$ 
49  repeat forever
50     $R_i := (i, s, true)$  /*register new name*/
51     $S := \{p_\ell \mid R_\ell \neq \perp\}$  /*collect suggested names*/
52    if  $\exists(\ell, s_\ell, b) \in S: i \neq \ell$  and  $s = s_\ell$  then
53       $r :=$  the rank of  $i$  in  $\{\ell \mid (\ell, s_\ell, b) \in S, b = true\}$ 
54       $s :=$  the  $r$ th integer not in  $\{s_\ell : (\ell, s_\ell, b) \in S, i \neq \ell\}$ 
55      /*rank among not yet decided participants*/
56      /*suggest a new name among not yet suggested*/
57    else
58       $R_i := (i, s, false)$ 
59    return  $s$ 

```

Figure 4: A k -concurrent $(j, j + k - 1)$ -renaming algorithm: code for each process p_i .

steps of \mathcal{A} in the resulting 2-concurrent run eventually decides and departs. The set of undecided participants gets smaller by one, and the next C -process with the smallest identity joins the 2-concurrent run of \mathcal{A} .

But it is shown in [15] that if all 1-resilient runs of a restricted algorithm \mathcal{A} satisfy strong j -renaming then there is a restricted algorithm to solve strong 2-renaming 2-concurrently—a contradiction with Lemma 11. \square

D.2 Solving renaming

The distributed algorithm used to solve $(j, j + k - 1)$ -renaming k -concurrently essentially mimics the algorithm of [3, 4] for wait-free $(j, 2j - 1)$ -renaming.

Theorem 15 *For all $1 < k \leq j < m$, $(j, j + k - 1)$ -renaming can be solved k -concurrently.*

Proof. Our algorithm, described in Figure 4, essentially mimics the algorithm of [3, 4] for wait-free $(j, 2j - 1)$ -renaming.

In the algorithm, every process periodically selects a new name according to the set of the names not yet suggested by other processes and its rank among the set of currently not yet decided participants (lines 53 and 54).

Note that since at most j processes participate in every run, p_i can observe at most $j - 1$ names suggested by other processes in line 51. Furthermore, since in a k -concurrent run, p_i can observe at most k not yet decided participants, its rank can be at most k . Therefore, the highest name p_i can suggest in line 50 is $j + k - 1$.

Now we show that no two processes output the same name. Suppose, by contradiction, that p_i and p_j output the same name s . Thus, both p_i and p_j previously suggested s in line 50. But since after that both processes read each other's registers, at least one of them would see that s has been suggested by another process and thus would not decide—a contradiction.

Finally, we show that every correct process eventually decides. Consider, by contradiction, an run R in which a set of correct processes $\{p_{j_1}, \dots, p_{j_t}\}$ (ordered by their ids) never decide. We call these processes *trying*. We establish a contradiction by showing that p_{j_1} must eventually decide. Indeed, consider R' , a prefix of R , in which only trying processes take steps, and let S be the set of names suggested by the processes not in $\{p_{j_1}, \dots, p_{j_t}\}$ (note that this set does not change in R). Since, p_{j_1} has the smallest rank among the trying processes (let us denote it by r), eventually no trying process will ever suggest the r th name not in S . Thus, p_{j_1} eventually finds

itself to be the only process to suggest the name and decides—a contradiction. □

From this result and Theorem 9, we can conclude:

Theorem 16 *For all $1 < k \leq j < m$, $(j, j + k - 1)$ -renaming can be solved with $\neg\Omega_k$.*